

(19)



Europäisches Patentamt  
European Patent Office  
Office européen des brevets



(11)

EP 0 772 363 A2

(12)

## EUROPEAN PATENT APPLICATION

(43) Date of publication:  
07.05.1997 Bulletin 1997/19

(51) Int. Cl.<sup>6</sup>: H04N 7/26, H04N 7/50,  
H04N 7/36

(21) Application number: 96306890.3

(22) Date of filing: 20.09.1996

(84) Designated Contracting States:  
DE FR

(30) Priority: 30.10.1995 GB 9522169

(71) Applicant: SONY UNITED KINGDOM LIMITED  
Weybridge KT13 0XW (GB)

(72) Inventors:  
• Saunders, Nicholas Ian  
Basingstoke, Hampshire RG24 0RL (GB)

• Gillard, Clive Henry  
Basingstoke, Hampshire RG22 5JA (GB)  
• Ludgate, Michael John  
Basingstoke, Hampshire RG21 3NJ (GB)  
• Stone, Jonathan James  
Reading, Berkshire RG7 3SS (GB)

(74) Representative: Turner, James Arthur et al  
D. Young & Co.,  
21 New Fetter Lane  
London EC4A 1DA (GB)

### (54) Video data compression

(57) Video data compression apparatus in which successive images are encoded, comprises a data quantity allocator for allocating respective quantities of output encoded data for two or more complementary areas of each image; and an encoder for encoding each

area of each image of the group of images in accordance with the quantity of output encoded data allocated to that area.

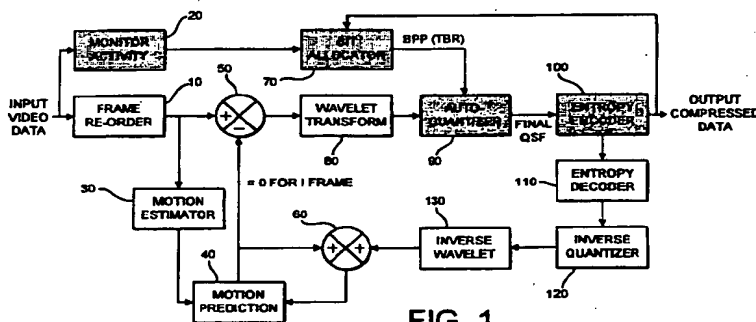


FIG. 1

## Description

This invention relates to video data compression.

Some video data compression systems, such as systems defined by the MPEG II standard<sup>1</sup>, use a number of different coding techniques to encode successive pictures of a video signal.

Typically, the video signal is divided into successive groups of pictures (GOPs). Within each GOP at least one picture is encoded as an "I-picture", or intra-picture, using only information present in that picture itself. This means that I-pictures can later be decoded without requiring information from other pictures, and so provide random entry points into the video sequence. However, the converse of this is that the encoding of I-pictures cannot make use of the similarity between successive pictures, and so the degree of data compression obtained with I-pictures is only moderate.

Further pictures within each GOP may be encoded as "P-pictures" or predicted pictures. P-pictures are encoded with respect to the nearest previous I-picture or P-picture, so that only the differences between a P-picture and the previous P- or I-picture needs to be transmitted. Also, motion compensation is used to encode the differences, so a much higher degree of compression is obtained than with I-pictures.

Finally, some of the pictures within a GOP may be encoded as "B-pictures" or bidirectional pictures. These are encoded with respect to two other pictures, namely the nearest previous I- or P-picture and the nearest following I- or P-picture. B-pictures are not used as references for encoding other pictures, so a still higher degree of compression can be used for B-pictures because any coding errors caused by the high compression will not be propagated to other pictures.

Therefore, in each GOP there are (up to) three classes of picture, I-, P- and B- pictures, which *tend to* achieve different degrees of compression and so *tend to* require different shares of the overall available encoded bit stream. Generally, I- pictures require a large share of the available transmission or storage capacity, followed by P-pictures, and followed by B-pictures.

In one previously proposed system, capacity (numbers of bits) in the encoded data stream is allocated between the various pictures in a GOP by a fixed picture to picture ratio. In another scheme, it has been proposed to allocate shares in the available encoded bit-stream to individual pictures within a GOP on the basis of the degree of compression obtained for corresponding pictures in a preceding GOP. In this way, the "I:B:P ratio", i.e. the ratio between the numbers of bits allocated to I-, B- and P-pictures in the encoded data stream can be steered in accordance with the image

content of the current pictures.

For example, in a sequence of pictures with little or no change between successive pictures, this steering method would result in the available data capacity being concentrated towards the I-pictures. At the opposite extreme, if there is a great deal of image motion between successive pictures, there is likely to be more difference data to be encoded for the B- and P- pictures and so the bit allocation technique would tend to allocate a greater share of the available capacity to B-and P- pictures.

It is a constant aim in the field of video compression to improve the subjective quality of the encoded (and subsequently decoded) images.

This invention provides video data compression apparatus in which successive images are encoded, the apparatus comprising:

- a data quantity allocator for allocating respective quantities of output encoded data for two or more complementary areas of each image; and
- an encoder for encoding each area of each image of the group of images in accordance with the quantity of output encoded data allocated to that area.

The invention recognises that the allocation of capacity in the encoded data stream merely on a picture-by-picture basis does not take into account the potentially differing encoding requirements of different areas of the *same picture*. Generally there will be different amounts of image motion and detail at different positions across the picture. An extreme example, however, would be when a cinema film is being broadcast in so-called "letterbox" format with black stripes at the top and bottom of the screen. Here, the coding requirements of the black stripes can be very different indeed to those of the film in the central portion of the screen.

The invention addresses this problem by performing bit allocation not only between whole pictures, but also between parts of a single picture. The parts could be, for example, 8x8 or 16x16 pixel macroblocks or even smaller, pairs of groups of macroblocks, horizontal or vertical stripes or slices of the picture, or other shapes. The general principle is that bit allocation can now take into account the differing needs of different parts of the picture.

The bit allocation could be fixed from GOP to GOP, or could be steered as described above. In the case of a fixed bit allocation, there is still scope for benefiting from allocating capacity between different areas of the picture, since (a) the viewer tends to notice coding errors more in the centre of the picture than at the edges, and (b) many domestic television sets overscan so that the extreme periphery of the picture is not seen by the home viewer at all. Accordingly, it is preferred that the allocation scheme weights the allocation towards the centre of the pictures. This weighting towards the centre is also of particular benefit with the "letterbox" transmission mentioned above.

<sup>1</sup> Motion Picture Experts Group II standard, described in the ISO/IEC Publication DIS 13818/2, "Information Technology - generic coding of motion pictures and associated audio information", March 1995

This technique is also applicable in a preferred embodiments in which successive GOPs are encoded subject to a predetermined limit on the quantity of output encoded data to be generated for each GOP, the apparatus comprising:

- a trial encoder for encoding at least a subset of images of a current group of images to generate trial-encoded data for each trial-encoded image;
- a data quantity allocator, responsive to the quantities of trial-encoded data generated by the trial encoder, for allocating respective quantities of output encoded data for each image in the group of images; and
- a final encoder for encoding each image of the group of images in accordance with the quantities of output encoded data allocated to that image by the data quantity allocator.

In the case of an allocation scheme (such as the steered scheme or the trial-encoding scheme mentioned above) which varies between GOPs, the bit allocation can now be made to reflect not only the "fixed" characteristics of the picture mentioned in the preceding paragraph but also the different image contents of parts of the pictures.

In order to manage situations such as a scene change or a change from "letterbox" to full screen transmission, it is preferred that each part of the picture receives at least a predetermined (non-zero) minimum bit allocation.

In a variable allocation scheme, the target bit rate allocation for each part of the picture is preferably made on the basis of the degree of detail or "activity" in that part of the picture. In fact, counter-intuitively, areas having a greater image activity (more detail) require a smaller bit allocation than areas having a low activity. This is because coding errors are subjectively more visible in areas of little detail.

The invention also provides a method of video data compression apparatus in which successive images are encoded, the method comprising the steps of:

- allocating respective quantities of output encoded data for two or more complementary areas of each image; and
- encoding each area of each image of the group of images in accordance with the quantity of output encoded data allocated to that area.

The invention will now be described by way of example with reference to the accompanying drawings, throughout which like parts are referred to by like references, and in which:

Figure 1 is a schematic diagram of a video data compression apparatus;

Figure 2 is a schematic diagram of a video picture illustrating different levels of image activity;

Figure 3 is a schematic diagram of an auto-quantiser;

Figure 4 is a schematic timing diagram illustrating the operation of the auto-quantiser of Figure 3;

Figure 5 is a schematic diagram of a second video data compression apparatus;

Figure 6 is a schematic diagram of a third video data compression apparatus;

Figure 7 is a schematic diagram of a fourth video data compression apparatus;

Figure 8 is a schematic diagram of a fifth video data compression apparatus; and

Figure 9 is a schematic diagram illustrating the weighting of target bit rates.

Figure 1 is a schematic diagram of a video data compression apparatus comprising a frame reorderer 10, an activity detector 20, a motion estimator 30, a motion predictor 40, a subtractor 50, an adder 60, a bit allocator 70, a wavelet transform unit 80, an auto-quantiser 90, an entropy encoder 100, an entropy decoder 110, an inverse quantiser 120 and an inverse wavelet coder 130.

(Where a reference is made in this description to wavelet encoding, it will be appreciated by the skilled man that other encoding techniques, for example the discrete cosine transform or DCT, may be used instead).

Many features of the apparatus of Figure 1 operate in a very similar manner to corresponding features of an MPEG encoder. Such features will not be described in detail here.

Briefly, therefore, the frame reorderer 10 receives input video data and acts on successive groups of pictures (GOP) to reorder the pictures so that each picture within the GOP is compressed *after* those pictures on which it depends. For example, if a B-picture (bi-directionally predicted picture) depends on a following I- or P-picture, it is reordered to be compressed after that I- or P-picture.

For example, if a GOP comprises the following four initial frames (in the order in which they are *displayed*),  $I_0 B_1 B_2 P_3 \dots$ , where the P-picture uses the I-picture as a reference and the two B-pictures use the surrounding I- and P-pictures as references, then the frame reorderer 10 will reorder the GOP to be *compressed* in the following order:  $I_0 P_3 B_1 B_2 \dots$

I-pictures are intra-picture encoded, that is to say the encoding is not based on any other reference pictures. An I-picture in a GOP is therefore passed from the frame reorderer 10 to the wavelet transform unit 80, the auto quantiser 90 and the entropy encoder 100 to generate output compressed data representing that I-picture.

The compressed I-picture data is also passed from the entropy encoder 100 through a decompression chain formed by the entropy decoder 110, the inverse quantiser 120, and the inverse wavelet transform unit 130. This reconstructs a version of the I-picture present

in the decoder which is passed to the motion predictor 40.

The next picture of the GOP to be compressed, which will generally be a P-picture which depends on the I- picture as a reference, is passed from the frame reorderer 10 to the motion estimator 30 which generates motion vectors indicative of image motion between the I- and P- pictures. The motion predictor 40 then generates a predicted version of the P picture using the motion vectors and the decoded version of the I- picture. This predicted version of the P- picture is subtracted from the actual P- picture by the subtractor 50 and the difference between the 2 frames is passed to the wavelet transform unit 80 for compression. As before, the encoded (compressed) difference data is output by the entropy encoder and is then decoded by the decompression chain 110, 120, 130 to regenerate a version of the difference data.

In the adder 60 the difference data is then added to the previously decompressed version of the I- picture to generate a decompressed version of the P-picture which is then stored in the motion predictor 40 for use in the compression of the next picture.

This process continues, so that each picture which uses other pictures as a reference is in fact compressed by encoding difference data between the input picture and a version of the input picture formed by motion prediction from a previously compressed and then decompressed version of the reference picture. This means that the compression is performed with respect to the pictures which will be available at the decompressor.

The activity detector 20 detects the image "activity", or "degree of detail" in blocks of each input image. This process will be described in more detail with reference to Figure 2 below.

The bit allocator 70 allocates target bit rates to whole pictures or blocks of the pictures in dependence on the image activity of pictures of the current GOP and the degree of quantisation obtained for I-, B and P- pictures of the preceding GOP. In fact, the allocation can be made by allocating an overall target bit rate for each GOP ( $TBR_{GOP}$ ) in proportions dependent on the actual quantity of data generated for the corresponding frame in the preceding GOP, or in accordance with the actual I:B:P ratio achieved with the preceding GOP. In this way, the allocation or the I:B:P ratio can be "steered" to reflect the type of image content in use.

The target bits rates are supplied to the auto quantiser 90 which generates a suitable quantisation factor to be applied to the wavelet encoded data to comply with the target bit rates. The auto quantiser will be described in more detail with reference to Figure 3 below.

Figure 2 is a schematic diagram of a video picture illustrating different levels of image activity.

The "activity" of a block of a picture is a measure of the degree of detail in that block. This means that a block representing a plain, bland area of the picture, such as an area of sky 190 in Figure 2 has a low image

activity. A block which includes a great deal of image detail such as an area 192 of the tree in Figure 1, has a high activity. Figure 1 also illustrates an intermediate area 194 at the edge of the tree, which has an activity value between that of the block 190 and that of the block 192.

The activity of each block of pixels in the image of Figure 2 can be measured by determining the variance ( $\sigma^2$ ) of the luminance values of the pixels in that block. The block activity is then derived by dividing the luminance variance for each block by the mean of variances for all of the blocks in the image. This yields a normalised block activity value:

$$\text{block activity} = \frac{\text{block } \sigma^2}{\text{mean } (\sigma^2)}$$

The image activity is monitored and sudden changes of at least a predetermined threshold between successive images of a GOP are detected. This type of sudden change can be indicative of a scene change occurring during the GOP. If such a sudden change is detected, then the bit allocator responds by resetting the "steered" I:B:P ratio to a default value intended to cope reasonably with the image content after the scene change.

A further justification for measuring image activity is that available space in a compressed video data stream can be allocated block-by-block (or by groups of blocks) depending on the respective measured block activities. In other words, the degree of data compression applied to each block or group of blocks is varied depending on the activity measured for that block or group of blocks.

The degree of compression can be varied (by varying the target bit rates supplied to the auto-quantiser 90) in dependence on the activity values because it has been observed that compression artifacts and distortions are more subjectively disturbing in blocks of low detail content (low activity, such as the plain sky area 190) than in blocks of high detail content (high activity, such as the block 192 of the tree in Figure 2). This is because the high level of detail in a block of high activity tends to distract the eye from the compression artifacts or distortions. Therefore, the target bit rates for areas or blocks of low image activity can be set to be higher than those for areas or blocks of high image activity.

One way of varying the target bit rates in response to the image activity is to increase or decrease the target bit rates by up to, say, 50%, in proportion to the image activity value. Here, the values should be normalised over the entire image or GOP so that the overall effect of the activity-value weighting of the target bit rates does not cause the target bit rates to exceed  $TBR_{GOP}$ .

Figure 3 is a schematic diagram of the auto-quantiser 90. (In fact, the apparatus of Figure 3 performs the function of the auto-quantiser 90 and the entropy encoder 100 of Figure 1).

The auto-quantiser 90 receives frequency separated video data on an input terminal 200 and an initial quantisation coefficient value,  $Q_s(0)$ , on an input terminal 210.

The initial  $Q_s$  value,  $Q_s(0)$ , may be a predetermined or preset number such as 64, or could be generated in response to the properties of the current video data being compressed.

The quantiser/encoder 240 quantises the current block of video data using the initial  $Q_s$  value supplied, and either performs a full Huffman coding process on the quantised data or simply detects the length of a variable length code which would result from the Huffman coding process (both of these are known techniques). In this way, the quantiser/encoder detects the amount of data which would result from quantising the input frequency separated video data using the supplied initial  $Q_s$  value.

A  $Q_s$  generator 232 compares the output bit rate of the quantiser/encoder 240 using the current  $Q_s$  value with the target bit rate for the current picture / area of a picture. If the detected output bit rate is lower than the target bit rate, then the quantisation scale factor should be reduced to apply less harsh quantisation to the frequency separated image data. Conversely, if the detected output bit rate is higher than the target bit rate, the quantisation scale factor should be increased. The  $q_s$  generator 232 therefore either increases or decreases the initial  $Q_s$  value,  $Q_s(0)$  by a predetermined proportion such as 50%, and supplies the modified  $Q_s$  value,  $Q_s(1)$  to a second quantiser/encoder 242. Accordingly, if the initial  $Q_s$  value is 4, then the value  $q_s(1)$  will be either 2 or 6.

This process is continued at each stage of the auto-quantiser, until the last stage is reached. There the  $q_s$  value which will result in the quantity of output data being close to but not exceeding the target bit rate for the current picture area is used as  $q_s(out)$ .

The apparatus of Figure 3 involves a three-stage binary search process to obtain the  $q_s$  value. It will be appreciated that in practice more stages could be used to define the  $q_s$  value to a higher resolution.

The final quantisation scale factor generated by this process,  $Q_s(out)$  may simply be output to a separate quantiser and entropy encoder. However, an alternative is illustrated in Figure 3, where a final quantiser/encoder 246 takes the place of a separate quantiser and the entropy encoder 100 of Figure 1, to generate the compressed output video.

The selection of the output quantisation scale factor  $Q_s(out)$  can be made for individual blocks of the image e.g.  $8 \times 8$  blocks of pixels, for groups of blocks, for regions of the image such as "slices" (horizontal strips of the image) or for an entire picture.

Figure 4 is a schematic timing diagram illustrating the operation of the auto-quantiser of Figure 3.

The examples shown in Figure 4 is that of processing each image as a series of "slices", ie horizontal stripes formed of groups of macroblocks. However, as

mentioned above, the staged processing illustrated in Figure 4 could be performed on various types of blocks of each picture. It is advantageous if the blocks on which the staged processing of Figure 4 is performed are co-sited with the blocks for which individual target bit rates are defined. In other words, co-sited blocks of the same size or larger than those for which the TBRs are defined are preferred.

The first process is that stage one of the auto quantiser (the quantiser/encoder 240 of Figure 3) processes a first slice of an image. When that processing is complete, that stage processors second, third, fourth and so on slices of the image.

When the processing of slice 1 has been completed by stage one, stage 2 takes on the processing of slice 1. When this has been completed, stage 3 works on slice 1, and finally when this has been completed, stage 4 works on slice 1. This technique leads to lower delays being caused by the auto-quantiser; if each stage had to wait until a preceding stage had completed processing of a whole image, the 4-stage auto-quantiser of Figure 3 would have a minimum delay of 4 picture periods. This delay is much reduced using the processing technique with individual blocks.

Figure 5 is a schematic diagram of a second video data compression apparatus. The apparatus of Figure 5 no longer uses the activity detector 20 or compression performance results from a previous GOP to set target bit rates for individual pictures of a GOP, but instead performs a trial or rehearsal compression and decompression of at least some of the pictures of each GOP.

The apparatus of Figure 5 may be considered as a trial encoder 300, which operates in much the same way as the apparatus of Figure 1, and a final encoder 310, which again operates in much the same way as the apparatus of Figure 1. The main difference between the trial encoder 300 and Figure 1 (apart from the lack of the activity detector and bit allocator) is that it uses a fixed quantiser 320 with a predetermined quantising factor.

In operation, the trial encoder 300 compresses a GOP (or at least some of the pictures of a GOP), using a fixed quantisation factor. The purpose of this is to assess the relative quantities of compressed data required by the different pictures or types of picture within the GOP.

While the trial compression is being performed, the GOP is stored in a delay unit 340.

The compressed data output by the entropy encoder of the trial compressor 300 is supplied to a bit allocator 330. The way in which the bit allocator 330 operates depends on whether all or some of the pictures are subject to trial compression.

If all of the pictures of a GOP are subject to trial compression, the bit allocator 330 then allocates target bit rates for final compression in accordance with the ratio of quantities of data generated during trial compression and subject to the overall limit on the storage or transmission channel's capacity (eg 30 Mbps - Mega

bits per second). In other words, each picture in a GOP is allocated a quantity of data according to the following derivation:

$$TBR_n = \frac{b_n}{\sum_0 b_n} \cdot TBR_{GOP}$$

where  $b_n$  is the quantity of trial data generated for the  $n$ th picture in the GOP and  $TBR_n$  is the target bit rate allocated to the  $n$ th picture.

If only some of the pictures in a GOP are subject to trial compression, but that sub set of pictures includes each type of picture present in the GOP, then an I:B:P ratio can be established for use with all pictures in that GOP. An I:B:P ratio, according to which the pictures in the GOP are then allocated quantities of compressed data, is derived using the following formulae:

$$I:B:P = \frac{\sum_{I\text{-pictures}} b_n}{N_{trial}(I)} : \frac{\sum_{P\text{-pictures}} b_n}{N_{trial}(P)} : \frac{\sum_{B\text{-pictures}} b_n}{N_{trial}(B)}$$

Finally, if the GOP contains I, B and P pictures, but only I- and P- pictures are subject to trial compression, a B:P ratio must be used to allocate quantities of data to the B pictures based on that available to I- and P- pictures. This B:P ratio could be fixed, or could be steered using the output of the final entropy encoder of the final compressor 310 having compressed the preceding GOP.

It will be appreciated that the above techniques can also be applied to sub-areas (e.g. blocks or "slices") of the pictures.

Once the bit allocation has been performed by the bit allocator 330, the GOP is passed from the delay unit 340 into the final compressor 310. The auto quantiser of the final compressor 310 controls the degree of quantisation applied to the image data so that the overall bit rate complies with the target bit rates generated by the bit allocator 330.

Figure 6 is a schematic diagram of a third video data compression apparatus.

Again, the apparatus of Figure 6 comprises a trial compressor 350 and a final compressor 310, with the trial compressor 350 supplying trial compressed data to a bit allocator 330 to control target bit rates to be used in final compression.

In contrast to the apparatus of Figure 5, the trial compressor 350 does not have a feedback loop of an entropy decoder, an inverse quantiser and an inverse wavelet coder. Instead, during trial compression, B- and P- pictures are compressed with respect to the input versions of the respective reference pictures. This means that the compression will be slightly less accu-

rate than a full compression using the feedback loop described above, but since the compression is only for the purposes of assessing approximate relative quantities of data required by each picture of the GOP, this slight inaccuracy does not matter and is more than compensated for by a saving in hardware or processing resources.

Figure 7 illustrates a similar arrangement being used in the trial compressor 350 and in a final compressor 360. Again, this can be useful where economy of processing resources is as important as compression accuracy.

Figure 8 is a schematic diagram of a fifth video data compression apparatus. This apparatus comprises a single compressor 400 which acts alternately as the trial compressor and the final compressor of any of the embodiments described earlier. This is achieved by running the processing of the compressor 400 at a faster speed than the data rate of the input video.

For example, a standard video data rate for the input video supplied to the frame reorderer 10 is 30.5 MHz. However, if the compressor 400 is operated at a basic video data rate of, say, 27 MHz, it has the ability to process the input video data of a GOP once for trial compression (in which the auto quantiser in fact operates at a fixed quantisation level) and then again for final compression in which the bit allocator 330 uses the quantities of data generated during trial compression to control the auto quantiser (now working with a variable quantisation factor).

If, for example, a 12 frame GOP is being processed and it is desired to perform trial quantisation on 4 frames of the GOP, the compressor 400 can be operated at 4/3 of the basic video data rate of 13.5 MHz, ie 18 MHz.

Finally, Figure 9 is a schematic diagram illustrating the weighting of target bit rates within an image.

As already described, target bit rates can be allocated to individual macroblocks, groups of macroblocks, slices, stripes and various other shaped blocks within an image. In the example shown in Figure 9, separate target bit rates are allocated to individual code blocks 410 of an image.

A weighting is applied to the target bit rates for blocks of the image, so that blocks towards the centre of the image are given higher target bit rates, all other factors being equal. This is done for two main reasons:

- (a) the viewer's centre of attention tends to be the middle of the picture, and so the viewer tends to notice coding errors more in the centre of the picture than at the edges, and
- (b) many domestic television sets overscan so that the extreme periphery of the picture is not seen by the home viewer at all.

Accordingly, the allocation scheme weights the allocation towards the centre of the pictures. This weighting towards the centre is also of particular benefit with so-called "letterbox" film transmission.

The graphs at the side and bottom of Figure 9 illustrate one possible envelope for weighting coefficients to be applied to the target bit rates for different blocks in the image.

In all of the above schemes for allocating target bit rates, each block, slice or image is allocated at least a predetermined minimum (non-zero) target bit rate by the bit allocator. This allows for image content not anticipated by the bit allocator and avoids areas of the picture being allocated such a small target bit rate that the image cannot be encoded.

In another preferred feature, the target bit rates (and in turn the degree of compression applied) can be varied in direct dependence on the picture type. In particular, in an I:B:P system a B picture (or areas of a B-picture) can be allocated a lower TBR than other pictures, since it is not used as a reference for any other pictures. This tends to lead to a higher bit allocation to pictures which are used as references for other pictures. However, although it may be expected that the subjective quality of the more harshly quantised images would then suffer, in fact, because the reference pictures are better, the non-reference pictures can also improve.

#### Claims

1. Video data compression apparatus in which successive images are encoded, the apparatus comprising:

a data quantity allocator for allocating respective quantities of output encoded data for two or more complementary areas of each image; and

an encoder for encoding each area of each image of the group of images in accordance with the quantity of output encoded data allocated to that area.

2. Apparatus according to claim 1, in which the allocation of output encoded data is in accordance with a predetermined allocation ratio between image areas.

3. Apparatus according to claim 1, in which the allocation of output encoded data between image areas for each image is dependent on the degree of image detail in each area of that image.

4. Apparatus according to any one of the preceding claims, in which the allocation of output encoded data between image areas is weighted to allocate a greater quantity of output encoded data to image areas towards a central region of the image.

5. Apparatus according to any one of the preceding claims, in which, in the allocation of output encoded data between image areas, each image area is allocated at least a predetermined minimum allocation

of output encoded data.

6. Apparatus according to any one of the preceding claims, in which successive groups of pictures are encoded subject to a predetermined limit on the quantity of output encoded data to be generated for each group of pictures, the apparatus comprising:

a trial encoder for encoding at least a subset of images of a current group of images to generate trial-encoded data for each trial-encoded image;

a data quantity allocator, responsive to the quantities of trial-encoded data generated by the trial encoder, for allocating respective quantities of output encoded data for each image in the group of images; and

a final encoder for encoding each image of the group of images in accordance with the quantities of output encoded data allocated to that image by the data quantity allocator.

7. A method of video data compression apparatus in which successive images are encoded, the method comprising the steps of:

allocating respective quantities of output encoded data for two or more complementary areas of each image; and  
encoding each area of each image of the group of images in accordance with the quantity of output encoded data allocated to that area.




FIG. 1







FIG. 2

FIG. 3




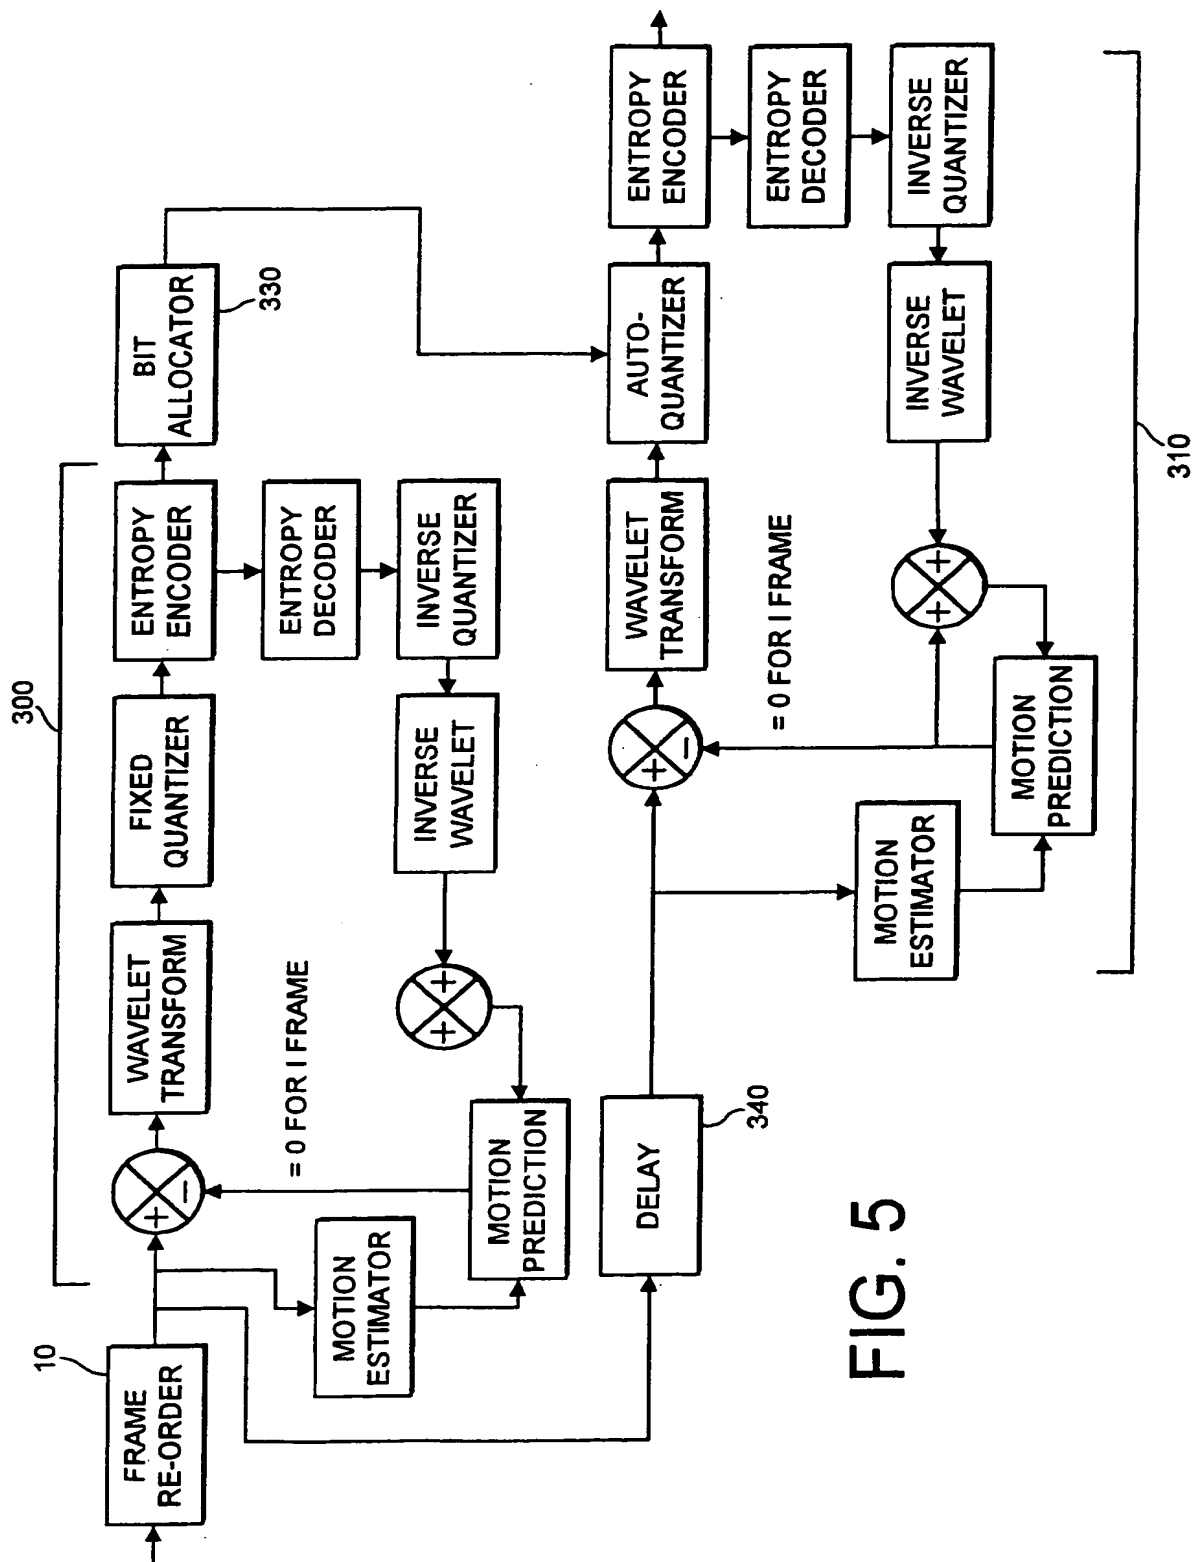



FIG. 4





6. G. F.




FIG. 7




FIG. 8




FIG. 9





(12)

**EUROPEAN PATENT APPLICATION**

(88) Date of publication A3:  
**21.07.1999 Bulletin 1999/29**

(51) Int. Cl.<sup>6</sup>: **H04N 7/26**, H04N 7/50,  
H04N 7/36, H04N 7/30

(43) Date of publication A2:  
**07.05.1997 Bulletin 1997/19**

**(21) Application number: 96306890.3**

**(22) Date of filing: 20.09.1996**

**(84) Designated Contracting States:**  
**DE FR**

**(30) Priority: 30.10.1995 GB 9522169**

**(71) Applicant:**  
**SONY UNITED KINGDOM LIMITED**  
**Weybridge KT13 0XW (GB)**

**(72) Inventors:**  
**• Saunders, Nicholas Ian**  
**Basingstoke, Hampshire RG24 ORL (GB)**

- **Gillard, Clive Henry**  
**Basingstoke, Hampshire RG22 5JA (GB)**
- **Ludgate, Michael John**  
**Basingstoke, Hampshire RG21 3NJ (GB)**
- **Stone, Jonathan James**  
**Reading, Berkshire RG7 3SS (GB)**

**(74) Representative:**  
**Turner, James Arthur et al**  
**D. Young & Co.,**  
**21 New Fetter Lane**  
**London EC4A 1DA (GB)**

**(54) Video data compression**

(57) Video data compression apparatus in which successive images are encoded, comprises a data quantity allocator for allocating respective quantities of output encoded data for two or more complementary areas of each image; and an encoder for encoding each

area of each image of the group of images in accordance with the quantity of output encoded data allocated to that area.




FIG. 1



European Patent  
Office

# EUROPEAN SEARCH REPORT

Application Number  
EP 96 30 6890

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
X	US 5 144 424 A (SAVATIER TRISTAN) 1 September 1992 * the whole document *	1-3,5-7	H04N7/26 H04N7/50 H04N7/36 H04N7/30
X	EP 0 541 302 A (AMERICAN TELEPHONE & TELEGRAPH) 12 May 1993 * page 3, line 32 - page 5, line 51 * * page 11, line 49 - page 11, line 57 *	1,3,5-7	
X	US 4 972 260 A (FUJIKAWA WATARU ET AL) 20 November 1990 * column 4, line 55 - column 5, line 12 * * figure 3 *	1,4,7	
X	US 5 402 146 A (RODRIGUEZ ARTURO A ET AL) 28 March 1995 * column 3, line 9 - column 5, line 17 *	1,4,7	
A	BAYLON D M ET AL: "VIDEO COMPRESSION AND NOISE REDUCTION USING TRANSFORM/SUBBAND CODING AND ADAPTIVE AMPLITUDE MODULATION" 1 June 1992, SMPTE JOURNAL, VOL. 101, NR. 6, PAGE(S) 404 - 406 XP000276817 * the whole document *	1-7	
The present search report has been drawn up for all claims			TECHNICAL FIELDS SEARCHED (Int.Cl.8)
			H04N
Place of search	Date of completion of the search	Examiner	
THE HAGUE	31 May 1999	Fassnacht, C	
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons &amp; : member of the same patent family, corresponding document</p>			

EPO FORM 1503 03.92 (P4/C01)

**ANNEX TO THE EUROPEAN SEARCH REPORT  
ON EUROPEAN PATENT APPLICATION NO.**

EP 96 30 6890

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.  
The members are as contained in the European Patent Office EDP file on  
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

31-05-1999

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5144424      A	01-09-1992	CN    1073313 A,B	16-06-1993
		DE    4233543 A	22-04-1993
		FR    2684833 A	11-06-1993
		GB    2260668 A,B	21-04-1993
		JP    5244586 A	21-09-1993
		MX    9205931 A	01-06-1993
EP 0541302      A	12-05-1993	US    5214507 A	25-05-1993
		CA    2081405 A	09-05-1993
		JP    2651097 B	10-09-1997
		JP    6070311 A	11-03-1994
US 4972260      A	20-11-1990	JP    2056187 A	26-02-1990
		JP    2056186 A	26-02-1990
US 5402146      A	28-03-1995	NONE	